

[0018] 1. An X-Ray machine M_1 that supported DICOM, would take a digital X-Ray of a patient. It would combine this X-Ray image with some additional information required by the DICOM standard such as how the image was taken, the specifications of the device etc. along with some limited information about the patient (such as his birth-date and provider-generated identifier). Much of this information will be generated by the X-Ray machine M_1 itself, with the operating technician providing all the data regarding the patient, body part scanned etc.

[0019] 2. The X-Ray image wrapped with the additional information “tags” so that it is now in DICOM format would be transferred across the network into the provider’s central PACS server 110 for storage. This means that the X-Ray machine M_1 needs limited local storage, backups etc. all of which would be provided at the PACS server 110 level, either within the PACS server system itself or by other systems controlled by and accessible to the PACS server 110 such as for example a group of computers that make up the PACS server 110’s local Storage Area Network.

[0020] 3. The PACS server 110 would deliver the obtained image to all hospital staff as needed making it available at Workstation WS_1 for example, with the X-Ray machine M_1 now free to work on other patients and images.

[0021] The important thing to note is that the X-Ray machine M_1 needs to have only the very limited intelligence needed to support the DICOM standard by wrapping its images with the right tags and being able to network with the PACS server 110. Such intelligence could be added at very low cost to virtually any imaging modality be it a CAT scan machine, an MRI scanner etc., a factor that was particularly important during the early 1990s when computing technology was much more expensive. The advancement of scanning technologies and the development of new modalities do not require any change to the provider workflow—as long as a new machine speaks DICOM, it will interact seamlessly with the central PACS server 110. Many U.S. hospitals have X-Ray machines that are often 20 or more years old because the actual scanning technology for capturing an X-Ray has not changed much over the period. If these machines are DICOM-aware, they still have all the necessary computing intelligence to capture their images and deliver them to the provider PACS server 110, however recent the vintage of the latter might be.

3.3 Limitations of the Prior Art DICOM/PACS Model

[0022] The DICOM standard was established at a time when high-speed networking was not widely available even within a single organization, the Internet did not exist and digital memory and storage were scarce and expensive. Most hospitals during this time worked with X-ray, CAT scan and other images in physical form. By creating a robust networking protocol, a centralized storage framework and consistent image tagging, DICOM allowed hospitals to replace their largely analog/physical workflow with digital images, networked imaging hardware and centralized storage. The DICOM standard was a huge step forward where it came to medical imaging and provided most of the functions that are needed within a single hospital for imaging.

[0023] In today’s world, computer networks are very fast and ubiquitous, networking protocols are robust and digital storage is exceptionally cheap and reliable. DICOM has adapted to some of these changes. Thus, for example, numerous web extensions for DICOM (DicomWeb) collectively allow for DICOM image handling over the Internet using standard Internet protocols. Again, DICOM networking speeds and PACS server storage amounts and image access speeds have all improved significantly with these advances in technology.

[0024] However, even with these improvements in implementation and the extensions to the standard engendered by technological change, the DICOM/PACS framework still remains a system that works best within a given medical provider. This is entirely understandable because DICOM was a standard originally devised to address the medical imaging needs of a single provider. Thus, even today, DICOM does not easily permit interoperability across medical providers except to a very limited degree. There are three major issues with the existing DICOM/PACS framework which we take up in the sections below.

3.3.1 Images Have Globally Unique Identifiers But Patients Do Not

[0025] The DICOM standard organizes images into a hierarchy of studies, series and instances with one study being made up of one or more series, and with the latter in turn being composed of one or more instances. The standard requires that each study, series and instance have a Globally Unique Identifier (UID) no matter what modality generated the image. Thus, a sequence of X-Ray images obtained on a machine in Hospital A for a specific patient would be grouped into studies, series and instances each of which would have a UID that cannot be repeated either in the same hospital or anywhere else. Were a second group of images to be taken, they would have to be identified with a new (and unique) group of study, series and instance identifiers even if they pertain to the same patient and the same body locations.

[0026] The DICOM standards setters have managed to ensure this global uniqueness even with the explosive growth in digital imaging with UIDs working reliably for the most part. Minor issues may arise within a hospital because of specific workflows. Thus, when a hospital edits the images generated by a DICOM modality using a particular software package, the latter has to take the proper care to ensure that these edited images are labeled with new study, series and instances UIDs as appropriate so that they do not conflict with the unique identifiers of the original image. More simply, a correct implementation of the rules established by the standard (which requires new UIDs for edited images) would ensure global uniqueness of the image UIDs in regard to study, series and instance.

[0027] There are other problems that might occur even when we have global uniqueness of study, series and instance UIDs. For example, we might have a patient at a provider who has 75 images taken during a single session that should logically be organized into 5 series (of say 15 images each) all consolidated into a single study. However, these images might be coded with each instance corresponding to a unique series and study, so that there are 75 studies and 75 series with each study and series corresponding to a single image instance. Such issues are not uncommon especially when the UID tags are not assigned by the modality